

What is Gab? A Bastion of Free Speech or an Alt-Right Echo Chamber?

Savvas Zannettou
Cyprus University of Technology
sa.zannettou@edu.cut.ac.cy

Barry Bradlyn
Princeton Center for Theoretical Science
bbradlyn@princeton.edu

Emiliano De Cristofaro
University College London
e.decrisofaro@ucl.ac.uk

Haewoon Kwak
Qatar Computing Research Institute
& Hamad Bin Khalifa University
hkwak@hbku.edu.qa

Michael Sirivianos
Cyprus University of Technology
michael.sirivianos@cut.ac.cy

Gianluca Stringhini
University College London
g.stringhini@ucl.ac.uk

Jeremy Blackburn
University of Alabama at Birmingham
blackburn@uab.edu

ABSTRACT

Over the past few years, a number of new “fringe” communities, like 4chan or certain subreddits, have gained traction on the Web at a rapid pace. However, more often than not, little is known about how they evolve or what kind of activities they attract, despite recent research has shown that they influence how false information reaches mainstream communities. This motivates the need to monitor these communities and analyze their impact on the Web’s information ecosystem.

In August 2016, a new social network called Gab was created as an alternative to Twitter. It positions itself as putting “people and free speech first”, welcoming users banned or suspended from other social networks. In this paper, we provide, to the best of our knowledge, the first characterization of Gab. We collect and analyze 22M posts produced by 336K users between August 2016 and January 2018, finding that Gab is predominantly used for the dissemination and discussion of news and world events, and that it attracts alt-right users, conspiracy theorists, and other trolls. We also measure the prevalence of hate speech on the platform, finding it to be much higher than Twitter, but lower than 4chan’s Politically Incorrect board.

CCS CONCEPTS

• **General and reference** → **General conference proceedings; Measurement; • Mathematics of computing** → **Probabilistic algorithms; • Networks** → **Social media networks; Online social networks;**

KEYWORDS

Social Networks, Hate Speech, Alt-Right, Changepoint Analysis

ACM Reference Format:

Savvas Zannettou, Barry Bradlyn, Emiliano De Cristofaro, Haewoon Kwak, Michael Sirivianos, Gianluca Stringhini, and Jeremy Blackburn. 2018. What is Gab? A Bastion of Free Speech or an Alt-Right Echo Chamber?. In *WWW ’18 Companion: The 2018 Web Conference Companion, April 23–27, 2018, Lyon, France*. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3184558.3191531>

1 INTRODUCTION

The Web’s information ecosystem is composed of multiple communities with varying influence [24]. As mainstream online social networks become less novel, users have begun to join smaller, more focused platforms. In particular, as the former have begun to reject fringe communities identified with racist and aggressive behavior, a number of alt-right focused services have been created. Among these emerging communities, the Gab social network has attracted the interest of a large number of users since its creation in 2016 [8], a few months before the US Presidential Election. Gab was created, ostensibly as a censorship-free platform, aiming to protect free speech above anything else. From the very beginning, site operators have welcomed users banned or suspended from platforms like Twitter for violating terms of service, often for abusive and/or hateful behavior. In fact, there is extensive anecdotal evidence that the platform has become the alt-right’s new hub [23] and that it exhibits a high volume of hate speech [13] and racism [5]. As a result, in 2017, both Google and Apple rejected Gab’s mobile apps from their stores because of hate speech [13] and non-compliance to pornographic content guidelines [1].

In this paper, we provide, to the best of our knowledge, the first characterization of the Gab social network. We crawl the Gab platform and acquire 22M posts by 336K users over a 1.5 year period (August 2016 to January 2018). Overall, the main findings of our analysis include:

- (1) Gab attracts a wide variety of users, ranging from well-known alt-right personalities like Milo Yiannopoulos to conspiracy theorists like Alex Jones. We also find a number of “troll” accounts that have migrated over from other platforms like 4chan, or that have been heavily inspired by them.

This paper is published under the Creative Commons Attribution 4.0 International (CC BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW ’18 Companion, April 23–27, 2018, Lyon, France

© 2018 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC BY 4.0 License.

ACM ISBN 978-1-4503-5640-4/18/04.

<https://doi.org/10.1145/3184558.3191531>

- (2) Gab is predominantly used for the dissemination and discussion of world events, news, as well as conspiracy theories. Interestingly, we note that Gab reacts strongly to events related to white nationalism and Donald Trump.
- (3) Hate speech is extensively present on the platform, as we find that 5.4% of the posts include hate words. This is 2.4 times higher than on Twitter, but 2.2 times lower than on 4chan’s Politically Incorrect board (/pol/) [9].
- (4) There are several accounts making coordinated efforts towards recruiting millennials to the alt-right.

In summary, our analysis highlights that Gab appears to be positioned at the border of mainstream social networks like Twitter and “fringe” Web communities like 4chan’s /pol/. We find that, while Gab claims to be all about free speech, this seems to be merely a shield behind which its alt-right users hide.

Paper Organization. In the next section, we review the related work. Then, in Section 3, we provide an overview of the Gab platform, while in Section 4 we present our analysis on Gab’s user base and the content that gets shared. Finally, the paper concludes in Section 5.

2 RELATED WORK

In this section, we review previous work on social network analysis and in particular on fringe communities.

Kwak et al. [11] are among the first to study Twitter, aiming to understand its role on the Web. They show that Twitter is a powerful network that can be exploited to assess human behavior on the Web. However, the Web’s information ecosystem does not naturally build on a single or a few Web communities; with this motivation in mind, Zannettou et al. [24] study how mainstream and alternative news propagate across multiple Web communities, measuring the influence that each community have on each other. Using a statistical model known as Hawkes Processes, they highlight that small “fringe” Web communities within Reddit and 4chan can have a substantial impact on large mainstream Web communities like Twitter.

With the same multi-platform point of view, Chandrasekharan et al. [6] propose an approach, called Bag of Communities, which aims to identify abusive content within a community. Using training data from nine communities within 4chan, Reddit, Voat, and Metafilter, they outperform approaches that focus only on in-community data.

Other work also focuses on characterizing relatively small alt-right Web communities. Specifically, Hine et al. [9] study 4chan’s Politically Incorrect board (/pol/), and show that it attracts a high volume of hate speech. They also find evidence of organized campaigns, called *raids*, that aim to disrupt the regular operation of other Web communities on the Web; e.g., they show how 4chan users raid YouTube videos by posting large numbers of abusive comments in a relatively small period of time.

Overall, our work is, to the best of our knowledge, the first to study the Gab social network, analyzing what kind of users it attracts, what are the main topics of discussions, and to what extent Gab users share hateful content.

3 GAB

Gab is a new social network, launched in August 2016, that “champions free speech, individual liberty, and the free flow of information online.”¹ It combines social networking features that exist in popular social platforms like Reddit and Twitter. A user can broadcast 300-character messages, called “gabs,” to their followers (akin to Twitter). From Reddit, Gab takes a modified voting system (which we discuss later). Gab allows the posting of pornographic and obscene content, as long as users label it as Not-Safe-For-Work (NSFW).² Posts can be reposted, quoted, and used as replies to other gabs. Similar to Twitter, Gab supports hashtags, which allow indexing and querying for gabs, as well as mentions, which allow users to refer to other users in their gabs.

Topics and Categories. Gab posts can be assigned to a specific *topic* or *category*. *Topics* focus on a particular event or timely topic of discussion and can be created by Gab users themselves; all topics are publicly available and other users can post gabs related to topics. *Categories* on the other hand, are defined by Gab itself, with 15 categories defined at the time of this writing. Note that assigning a gab to a category and/or topic is *optional*, and Gab moderates topics, removing any that do not comply with the platform’s guidelines.

Voting system. Gab posts can get up- and down-voted; a feature that determines the popularity of the content in the platform (akin to Reddit). Additionally, each user has its own score, which is the sum of up-votes minus the sum of down-votes that it received to all his posts (similar to Reddit’s user karma score [3]). This user-level score determines the popularity of the user and is used in a way unique to Gab: a user must have a score of at least 250 points to be able to down-vote other users’ content, and every time a user down-votes a post a point from his user-level score is deducted. In other words, a user’s score is used as a form of currency expended to down-vote content.

Moderation. Gab has a lax moderation policy that allows most things to be posted, with a few exceptions. Specifically, it only forbids posts that contain “illegal pornography” (legal pornography is permitted), posts that promote terrorist acts, threats to other users, and doxing other users’ personal information [18].³

Monetization. Gab is ad-free and relies on direct user support. On October 4, 2016 Gab’s CEO Andrew Torba announced that users were able to donate to Gab [19]. Later, Gab added “pro” accounts as well. “Pro” users pay a per-month fee granting additional features like live-stream broadcasts, account verification, extended character count (up to 3K characters per gab), special formatting in posts (e.g., italics, bold, etc.), as well as premium content creation. The latter allows users to create “premium” content that can only be seen by subscribers of the user, which are users that pay a monthly fee to the content creator to be able to view his posts. The premium content model allows for crowdfunding particular Gab users, similar to the way that Twitch and Patreon work. Finally, Gab is in the process of raising money through an Initial Coin Offering (ICO) with the goal to offer a “censorship-proof” peer-to-peer social network that developers can build application on top [2].

¹<http://gab.ai>

²What constitutes NSFW material is not well defined.

³For more information on Gab’s guidelines, see <https://gab.ai/about/guidelines>.

Followers			Scores			PageRank		
Name	Username	#	Name	Username	#	Name	Username	PR score
Milo Yiannopoulos	m	45,060	Andrew Torba	a	819,363	Milo Yiannopoulos	m	0.013655
PrisonPlanet	PrisonPlanet	45,059	John Rivers	JohnRivers	606,623	Andrew Torba	a	0.012818
Andrew Torba	a	38,101	Ricky Vaughn	Ricky_Vaughn99	496,962	PrisonPlanet	PrisonPlanet	0.011762
Ricky Vaughn	Ricky_Vaughn99	30,870	Don	Don	368,698	Mike Cernovich	Cernovich	0.006549
Mike Cernovich	Cernovich	29,081	Jared Wyand	JaredWyand	281,798	Ricky Vaughn	Ricky_Vaughn99	0.006143
Stefan Molyneux	stefanmolyneux	26,337	[omitted]	TukkRivers	253,781	Sargon of Akkad	Sargonofakkad100	0.005823
Brittany Pettibone	BrittPettibone	24,799	Brittany Pettibone	BrittPettibone	244,025	[omitted]	d_seaman	0.005104
Jebs	DeadNotSleeping	22,659	Tony Jackson	USMC-Devildog	228,370	Stefan Molyneux	stefanmolyneux	0.004830
[omitted]	TexasYankee4	20,079	[omitted]	causticbob	228,316	Brittany Pettibone	BrittPettibone	0.004218
[omitted]	RightSmarts	20,042	Constitutional Drunk	USSANews	224,261	Vox Day	voxday	0.003972
Vox Day	voxday	19,454	Truth Whisper	truthwhisper	206,516	Alex Jones	RealAlexJones	0.003345
[omitted]	d_seaman	18,080	Andrew Anglin	AndrewAnglin	203,437	Lauren Southern	LaurenSouthern	0.002984
Alex Jones	RealAlexJones	17,613	Kek_Magician	Kek_Magician	193,819	Donald J Trump	realdonaldtrump	0.002895
Jared Wyand	JaredWyand	16,975	[omitted]	shorty	169,167	Dave Cullen	DaveCullen	0.002824
Ann Coulter	AnnCoulter	16,605	[omitted]	SergeiDimitrovicIvanov	169,091	[omitted]	e	0.002648
Lift	lift	16,544	Kolja Bonke	KoljaBonke	160,246	Chuck C Johnson	Chuckcjohnson	0.002599
Survivor Medic	SurvivorMed	16,382	Party On Weimerica	CuckShamer	155,021	Andrew Anglin	AndrewAnglin	0.002599
[omitted]	SalguodNos	16,124	PrisonPlanet	PrisonPlanet	154,829	Jared Wyand	JaredWyand	0.002504
Proud Deplorable	luther	15,036	Vox Day	voxday	150,930	Pax Dickinson	pax	0.002400
Lauren Southern	LaurenSouthern	14,827	W.O. Cassity	wocassity	144,875	Baked Alaska	apple	0.002292

Table 1: Top 20 popular users on Gab according to the number of followers, their score, and their ranking based on PageRank in the followers/followings network. We omit the “screen names” of certain accounts for ethical reasons.

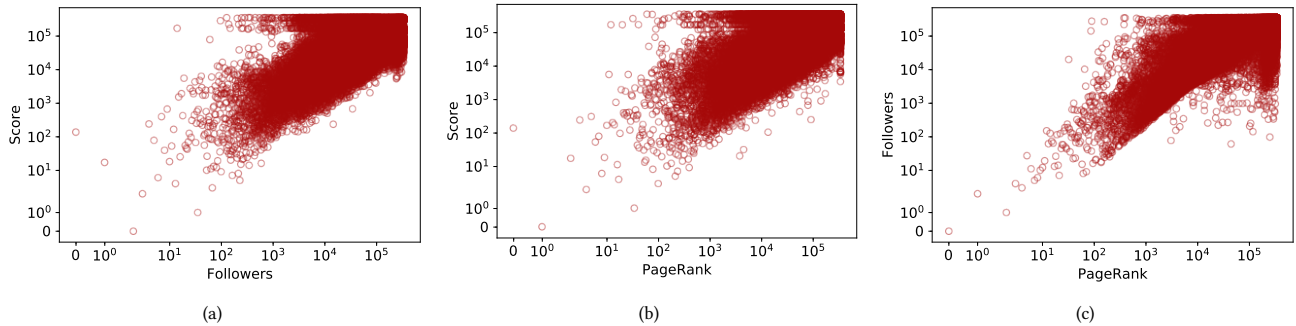


Figure 1: Correlation of the rankings for each pair of rankings: (a) Followers - Score; (b) PageRank - Score; and (c) PageRank - Followers.

Dataset. Using Gab’s API, we crawl the social network using a snowball methodology. Specifically, we obtain data for the most popular users as returned by Gab’s API and iteratively collect data from all their followers as well as their followings. We collect three types of information: 1) basic details about Gab accounts, including username, score, date of account creation; 2) all the posts for each Gab user in our dataset; and 3) all the followers and followings of each user that allow us to build the following/followers network. Overall, we collect 22,112,812 posts from 336,752 users, between August 2016 and January 2018.

4 ANALYSIS

In this section, we provide our analysis on the Gab platform. Specifically, we analyze Gab’s user base and posts that get shared across several axes.

4.1 Ranking of users

To get a better handle on the interests of Gab users, we first examine the most popular users using three metrics: 1) the number of followers; 2) user account score; and 3) user PageRank. These three metrics provide us a good overview of things in terms of “reach,” appreciation of content production, and importance in terms of position within the social network. We report the top 20 users for each metric in Table 1. Although we believe that their existence in Table 1 is arguably indicative of their public figure status, for ethical reasons, we omit the “screen names” for accounts in cases where a potential link between the screen name and the user’s real life names existed *and* it was unclear to us whether or not the user is a public figure. While Twitter has many celebrities in the most popular users [11], Gab seems to have what can at best be described as alt-right celebrities like Milo Yiannopoulos and Mike Cernovich.

Word	(%)	Bigram	(%)
maga	4.35%	free speech	1.24%
twitter	3.62%	trump supporter	0.74%
trump	3.53%	night area	0.49%
conservative	3.47%	area wanna	0.48%
free	3.08%	husband father	0.45%
love	3.03%	check link	0.42%
people	2.76%	freedom speech	0.41%
life	2.70%	hey guys	0.40%
like	2.67%	donald trump	0.40%
man	2.49%	man right	0.39%
truth	2.46%	america great	0.39%
god	2.45%	link contracts	0.35%
world	2.44%	wanna check	0.34%
freedom	2.29%	make america	0.34%
right	2.27%	need man	0.34%
american	2.25%	guys need	0.33%
want	2.23%	president trump	0.32%
one	2.20%	guy sex	0.31%
christian	2.17%	click link	0.30%
time	2.14%	link login	0.30%

Table 2: Top 20 words and bigrams found in the descriptions of Gab users.

Number of followers. The number of followers that each account has can be regarded as a metric of impact on the platform, as a user with many followers can share its posts to a large number of other users. We observe a wide variety of different users; 1) popular alt-right users like Milo Yiannopoulos, Mike Cernovich, Stefan Molyneux, and Brittany Pettibone; 2) Gab’s founder Andrew Torba; and 3) popular conspiracy theorists like Alex Jones. Notably lacking are users we might consider as counter-points to the alt-right right, an indication of Gab’s heavily right-skewed user-base.

Score. The score of each account is a metric of content popularity, as it determines the number of up-votes and down-votes that they receive from other users. In other words, is the degree of appreciation from other users. By looking at the ranking using the score, we observe two new additional categories of users: 1) users purporting to be news outlets, likely pushing false or controversial information on the network like PrisonPlanet and USSANews; and 2) troll users that seem to have migrated from or been inspired by other platforms (e.g., 4chan) like Kek_Magician and CuckShamer.

PageRank. We also compute PageRank on the followers/followings network and we rank the users according to the obtained score. We use this metric as it quantifies the structural importance of nodes within a network according to its connections. Here, we observe some interesting differences from the other two rankings. For example, the account with username “realdonaldtrump,” an account reserved for Donald Trump, appears in the top users mainly because of the extremely high number of users that follow this account, despite the fact that it has no posts or score.

Comparison of rankings. To compare the three aforementioned rankings, we plot the ranking of all the users for each pair of rankings in Fig. 1. We observe that the pair with the most agreement

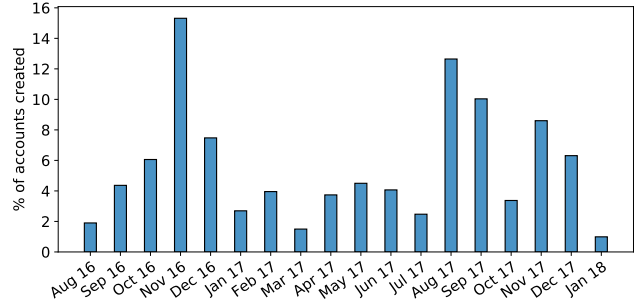


Figure 2: Percentage of accounts created per month.

is PageRank-Followers (Fig. 1(c)), followed by the pair Followers-Score (Fig. 1(a)), while the pair with the least agreement is PageRank - Score (Fig 1(b)). Overall, for all pairs we find a varying degree of rank correlation. Specifically, we calculate the Spearman’s correlation coefficient for each pair of rankings; finding 0.53, 0.42, 0.26 for PageRank-Followers, Followers-Score, and PageRank-Score, respectively. While these correlations are not terribly strong, they are significant ($p < 0.01$) for the two general classes of users: those that play an important structural role in the network, perhaps encouraging the diffusion of information, and those that produce content the community finds valuable.

4.2 User account analysis

User descriptions. To further assess the type of users that the platform attracts we analyze the description of each created account in our dataset. Note that by default Gab adds a quote from a famous person as the description of each account and a user can later change it. Although not perfect, we look for any user description enclosed in quotes with a “-” followed by a name, and assume it is a default quote. Using this heuristic, we find that only 20% of the users actively change their description from the default. Table 2 reports the top words and bigrams found in customized descriptions (we remove stop words for more meaningful results). Examining the list, it is apparent that Gab users are conservative Americans, religious, and supporters of Donald Trump and “free speech.” We also find some accounts that are likely bots and trying to deceive users with their descriptions; among the top bigrams there some that nudge users to click on URLs, possibly malicious, with the promise that they will get sex. For example, we find many descriptions similar to the following: “Do you wanna get sex tonight? One step is left ! Click the link - < url >.” It is also worth noting that our account (created for crawling the platform) was followed by 12 suspected bot accounts between December 2017 and January 2018 without making any interactions with the platform (i.e., our account has never made a post or followed any user).

User account creation. We also look when users joined the Gab platform. Fig. 2 reports the percentage of accounts created for each month of our dataset. Interestingly, we observe that we have peaks for account creation on November 2016 and August 2017. These findings highlight the fact that Gab became popular during notable world and politics events like the 2016 US elections as well as the Charlottesville Unite the Right rally [22]. Finally, only a small

percentage of Gab’s users are either pro or verified, 0.75% and 0.5%, respectively, while 1.7% of the users have a private account (i.e., only their followers can see their gabs).

Followers/Followings. Fig. 3 reports our analysis based on the number of followers and followings for each user. From Fig. 3(a) we observe that in general Gab users have a larger number of followers when compared with following users. Interestingly, 43% of users are following zero other users, while only 4% of users have zero followers. I.e., although counter-intuitive, most users have more followers than users they follow. Figs. 3(b) and 3(c) show the number of followers and following in conjunction with the number of posts for each Gab user. We bin the data in log-scale bins and we report the mean and median value for each bin. We observe that in both cases, that there is a near linear relationship with the number of posts and followers/followings up until around 10 followers/followings. After this point, we see this relationship diverge, with a substantial number of users with huge numbers of posts, some over 77K. This demonstrates the extremely heavy tail in terms of content production on Gab, as is typical of most social medial platforms.

Reciprocity. From the followers/followings network we find a low level of reciprocity: specifically, only 29.5% of the node pairs in the network are connected both ways, while the remaining 71.5% are connected one way. When compared with the corresponding metric on Twitter [11], these results highlight that Gab has a larger degree of network reciprocity indicating that the community is more tightly-knit, which is expected when considering that Gab mostly attracts users from the same ideology (i.e., alt-right community).

4.3 Posts Analysis

Basic Statistics. First, we note that 63% of the posts in our dataset are original posts while 37% are reposts. Interestingly, only 0.14% of the posts are marked as NSFW. This is surprising given the fact that one of the reasons that Apple rejected Gab’s mobile app is due to the share of NSFW content [1]. From browsing the Gab platform, we also can anecdotally confirm the existence of NSFW posts that are not marked as such, raising questions about how Gab moderates and enforces the use of NSFW tags by users. When looking a bit closer at their policies, Gab notes that they use a 1964 United States Supreme Court Ruling [21] on pornography that provides the famous “I’ll known it when I see it” test. In any case, it would seem that Gab’s social norms are relatively lenient with respect to what is considered NSFW.

We also look into the languages of the posts, as returned by Gab’s API. We find that Gab’s API does not return a language code for 56% of posts. By looking at the dataset, we find that all posts before June 2016 do not have an associated language; possibly indicating that Gab added the language field afterwards. Nevertheless, we find that the most popular languages are English (40%), Deutsch (3.3%), and French (0.14%); possibly shedding light to Gab’s users locations which are mainly the US, the UK, and Germany.

URLs. Next, we assess the use of URLs in Gab; overall we find 3.5M unique URLs from 81K domains. Table 3 reports the top 20 domains according to their percentage of inclusion in all posts. We observe that the most popular domain is YouTube with almost 7%

Domain	(%)	Domain	(%)
youtube.com	4.22%	zerohedge.com	0.53%
youtu.be	2.67%	twimg.com	0.53%
twitter.com	1.96%	dailycaller.com	0.49%
breitbart.com	1.44%	t.co	0.47%
bit.ly	0.82%	ussanews.com	0.46%
thegatewaypundit.com	0.74%	dailymail.co.uk	0.46%
kek.gg	0.69%	tinyurl.com	0.44%
imgur.com	0.68%	wordpress.com	0.43%
sli.mg	0.61%	foxnews.com	0.41%
infowars.com	0.56%	blogspot.com	0.32%

Table 3: Top 20 domains in posts and their respective percentage over all posts.

Hashtag	(%)	Mention	(%)
MAGA	6.06%	a	0.69%
GabFam	4.22%	TexasYankee4	0.31%
Trump	3.01%	Stargirlx	0.26%
SpeakFreely	2.28%	YouTube	0.24%
News	2.00%	support	0.23%
Gab	0.88%	Amy	0.22%
DrainTheSwamp	0.71%	RaviCrux	0.20%
AltRight	0.61%	u	0.19%
Pizzagate	0.57%	BlueGood	0.18%
Politics	0.53%	HorrorQueen	0.17%
PresidentTrump	0.47%	Sockalexis	0.17%
FakeNews	0.41%	Don	0.17%
BritFam	0.37%	BrittPettibone	0.16%
2A	0.35%	TukkRivers	0.15%
maga	0.32%	CurryPanda	0.15%
NewGabber	0.28%	Gee	0.15%
CanFam	0.27%	e	0.14%
BanIslam	0.25%	careyetta	0.14%
MSM	0.22%	PrisonPlanet	0.14%
1A	0.21%	JoshC	0.12%

Table 4: Top 20 hashtags and mentions found in Gab. We report their percentage over all posts.

of all posts, followed by Twitter with 2%. Interestingly, we note the extensive use of alternative news sources like Breitbart (1.4%), The Gateway Pundit (0.7%), and Infowars (0.5%), while mainstream news outlets like Fox News (0.4%) and Daily Mail (0.4%) are further below. Also, we note the use of image hosting services like Imgur (0.6%), sli.mg (0.6%), and kek.gg (0.7%) and URL shorteners like bit.ly (0.8%) and tinyurl.com (0.4%). Finally, it is worth mentioning that The Daily Stormer, a well known neo-Nazi web community is five ranks ahead of the most popular mainstream news source, The Hill.

Hashtags & Mentions As discussed in Section 3, Gab supports the use of hashtags and mentions similar to Twitter. Table 4 reports the top 20 hashtags/mentions that we find in our dataset. We observe that the majority of the hashtags are used in posts about Trump, news, and politics. We note that among the top hashtags are “AltRight”, indicating that Gab users are followers of the alt-right movement or they discuss topics related to the alt-right; “Pizzagate”, which denotes discussions around the notorious conspiracy theory [20]; and “BanIslam”, which indicate that Gab users are

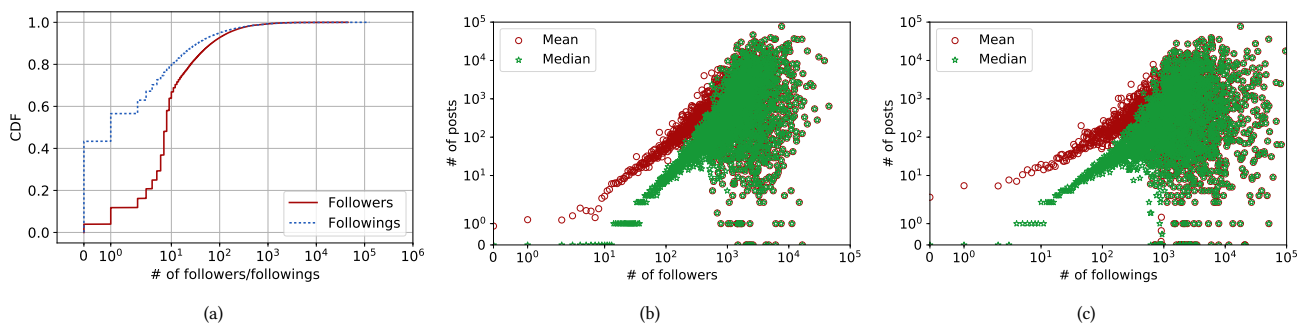


Figure 3: Followers and Following analysis (a) CDF of number of followers and following (b) number of followers and number of posts and (c) number of following and number of posts.

sharing their islamophobic views. It is also worth noting the use of hashtags for the dissemination of popular memes, like the Drain the Swamp meme that is popular among Trump’s supporters [14]. When looking at the most popular users that get mentioned, we find popular users related to the Gab platform like Andrew Torba (Gab’s CEO with username @a).

We also note users that are popular with respect to mentions, but do *not* appear in Table 1’s lists of popular users. For example, Amy is an account purporting to be Andrew Torba’s mother. The user Stargirlx, who we note changed usernames three times during our collection period, appears to be an account presenting itself as a millennial “GenZ” young woman. Interestingly, it seems that Amy and Stargirlx have been organizing Gab “chats,” which are private groups of users, for 18 to 29 year olds to discuss politics; possibly indicating efforts to recruit millennials to the alt-right community.

Categories & Topics. As discussed in Section 3 gabs may be part of a topic or category. By analyzing the data, we find that this happens for 12% and 42% of the posts for topics and categories, respectively. Table 5 reports the percentage of posts for each category as well as for the top 15 topics. For topics, we observe that the most popular are general “Ask Me Anything” (AMA) topics like Deutsch (2.29%, for German users), BritFam (0.73%, for British users), and Introduce Yourself (0.59%). Furthermore, other popular topics include world events and news like International News (0.59%), Las Vegas shooting (0.27%), and conspiracy theories like Seth Rich’s Murder (0.11%). When looking at the top categories we find that by far the most popular categories are News (15.91%) and Politics (10.30%). Other popular categories include AMA (4.46%), Humor (3.50%), and Technology (1.44%).

These findings highlight that Gab is heavily used for the dissemination and discussion of world events and news. Therefore, its role and influence on the Web’s information ecosystem should be assessed in the near future. Also, this categorization of posts can be of great importance for the research community as it provides labeled ground truth about discussions around a particular topic and category.

Hate speech assessment. As previously discussed, Gab was openly accused of allowing the dissemination of hate speech. In fact, Google removed Gab’s mobile app from its Play Store because it violates their hate speech policy [13]. Due to this, we aim to assess the extent of hate speech in our dataset. Using the modified Hatebase [4]

Topic	(%)	Category	(%)
Deutsch	2.29%	News	15.91%
BritFam	0.73%	Politics	10.30%
Introduce Yourself	0.59%	AMA	4.46%
International News	0.19%	Humor	3.50%
DACA	0.17%	Technology	1.44%
Las Vegas Terror Attack	0.16%	Philosophy	1.06%
Hurricane Harvey	0.16%	Entertainment	1.01%
Gab Polls	0.13%	Art	0.72%
London	0.12%	Faith	0.69%
2017 Meme Year in Review	0.12%	Science	0.56%
Twitter Purge	0.12%	Music	0.52%
Seth Rich	0.11%	Sports	0.39%
Memes	0.11%	Photography	0.37%
Vegas Shooting	0.11%	Finance	0.31%
Judge Roy Moore	0.09%	Cuisine	0.16%

Table 5: Top 15 categories and topics found in the Gab dataset

dictionary used by the authors of [9], we find that 5.4% of all Gab posts include a hate word. In comparison, Gab has 2.4 times the rate of hate words when compared to Twitter, but less than half the rate of hate words compared to 4chan’s Politically Incorrect board (/pol/) [9]. These findings indicate that Gab resides on the border of mainstream social networks like Twitter and fringe Web communities like 4chan’s Politically Incorrect (/pol/) board.

Temporal Analysis. Finally, we study the posting behavior of Gab users from a temporal point of view. Fig. 4 shows the distribution of the Gab posts in our dataset according to each day of our dataset, as well as per hour of day and week (in UTC). We observe that the general trend is that the number of Gab’s posts increase over time (Fig. 4(a)); this indicates an increase in Gab’s popularity. Furthermore, we note that Gab users posts most of their gabs during the afternoon and late night (after 3 PM UTC) while they rarely post during the morning hours (Fig. 4(b)). Also, the aforementioned posting behavior follow a diurnal weekly pattern as we show in Fig. 4(c).

To isolate significant days in the time series in Fig. 4(a), we perform a *change point analysis* using the Pruned Exact Linear Time (PELT) method [10]. First, we use our knowledge of the weekly variation in average post numbers from Fig. 4(c) to subtract from

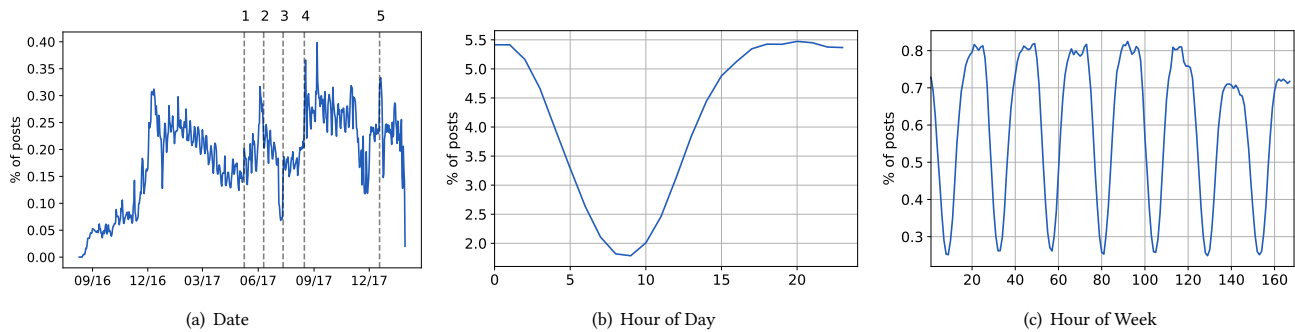


Figure 4: Temporal analysis of the Gab posts (a) each day; (b) based on hour of day and (c) based on hour of week.

our timeseries the mean number of posts for each day. This leaves us with a mean-zero timeseries of the deviation of the number of posts per day from the daily average. We assume that this timeseries is drawn from a normal distribution, with mean and variance that can change at a discrete number of changepoints. We then use the PELT algorithm to maximize the log-likelihood function for the mean(s) and variance(s) of this distribution, with a penalty for the number of changepoints. By ramping down the penalty function, we produce a ranking of the changepoints.

Examining current events around these changepoints provides insight into their dynamics that drive Gab behavior. First, we note that there is a general increase in activity up to the Trump inauguration, at which point activity begins to decline. When looking later down the timeline, we see an increase in activity after the changepoint marked 1 in Fig. 4(a). Changepoint 1 coincides with James Comey’s firing from the FBI, and the relative acceleration of the Trump-Russian collusion probe [16].

The next changepoint (2) coincides with the so-called “March Against Sharia” [12] organized by the alt-right, with the event marked 4 corresponding to Trump’s “blame on both sides” response to violence at the Unite the Right Rally in Charlottesville [15]. Similarly, we see a meaningful response to Twitter’s banning of abusive users [7] marked as changepoint 5.

Changepoint 3, occurring on July 12, 2017 is of particular interest, since it is the most extreme reduction in activity recognized as a changepoint. From what we can tell, this is a reaction to Donald Trump Jr. releasing emails that seemingly evidenced his meeting with a Russian lawyer to receive compromising intelligence on Hillary Clinton’s campaign [17]. I.e., the disclosure of evidence of collusion with Russia corresponded to the single largest drop in posting activity on Gab.

5 CONCLUSION

In this work, we have provided the first characterization of a new social network called Gab. We analyzed 22M posts from 336K users, finding that Gab attracts the interest of users ranging from alt-right supporters and conspiracy theorists to trolls. We showed that Gab is extensively used for the discussion of news, world events, and politics-related topics, further motivating the need to take it into account when studying information cascades on the Web. By looking at the posts for hate words, we also found that 5.4% of the posts include hate words. Finally, using changepoint analysis,

we highlighted how Gab reacts very strongly to real-world events focused around white nationalism and support of Donald Trump.

Acknowledgments. This project has received funding from the European Union’s Horizon 2020 Research and Innovation program under the Marie Skłodowska-Curie ENCASE project (Grant Agreement No. 691025). The work reflects only the authors’ views; the Agency and the Commission are not responsible for any use that may be made of the information it contains.

REFERENCES

- [1] 2016. Apple’s Double Standards Against Gab. <https://medium.com/@getongab/apples-double-standards-against-gab-1bffa2c09115>. (2016).
- [2] 2017. A Censorship-Proof P2P Social Media Protocol. <https://www.startengine.com/gab-select>. (2017).
- [3] 2017. Reddit FAQ - Karma. <https://www.reddit.com/wiki/faq>. (2017).
- [4] 2018. Hatebase API. <https://www.hatebase.org/>. (2018).
- [5] Thor Benson. 2016. Inside the “Twitter for racists”: Gab – the site where Milo Yiannopoulos goes to troll now. <https://goo.gl/Yqv4Ue>. (2016).
- [6] Eshwar Chandrasekharan, Mattia Samory, Anirudh Srinivasan, and Eric Gilbert. 2017. The Bag of Communities: Identifying Abusive Behavior Online with Preexisting Internet Data. In *CHI*.
- [7] Kitty Donaldson and Joshua Brustein. 2017. Twitter Bans Some White Supremacists and Other Extremists. <https://www.bloomberg.com/news/articles/2017-12-19/twitter-bans-some-white-supremacists-and-other-extremists>. (2017).
- [8] Gab. 2017. Gab site. <https://gab.ai>. (2017).
- [9] Gabriel Emile Hine, Jeremiah Onalapo, Emiliano De Cristofaro, Nicolas Kourtellis, Ilias Leontiadis, Riginos Samaras, Gianluca Stringhini, and Jeremy Blackburn. 2017. Kek, Cucks, and God Emperor Trump: A Measurement Study of 4chan’s Politically Incorrect Forum and Its Effects on the Web. In *ICWSM*.
- [10] R. Killick, P. Fearnhead, and I. A. Eckley. 2012. Optimal Detection of Changepoints With a Linear Computational Cost. *J. Amer. Statist. Assoc.* 107, 500 (2012), 1590–1598. <https://doi.org/10.1080/01621459.2012.737745> arXiv:<https://doi.org/10.1080/01621459.2012.737745>
- [11] Haewoon Kwak, Changhyun Lee, Hosung Park, and Sue B. Moon. 2010. What is Twitter, a social network or a news media?. In *WWW*.
- [12] Christopher Mathias. 2017. The ‘March Against Sharia’ Protests Are Really Marches Against Muslims. https://www.huffingtonpost.com/entry/march-against-sharia-anti-muslim-act-for-america_us_5939576ee4b0b13f2c67d50c. (2017).
- [13] Rob Price. 2017. Google’s app store has banned Gab, a social network popular with the far-right, for ‘hate speech’. <http://uk.businessinsider.com/google-app-store-gab-ban-hate-speech-2017-8>. (2017).
- [14] Aja Romano. 2016. The 2016 culture war, as illustrated by the alt-right. <https://www.vox.com/culture/2016/12/30/13572256/2016-trump-culture-war-alt-right-meme>. (2016).
- [15] Michael Shear and Maggie Haberman. 2017. Trump Defends Initial Remarks on Charlottesville; Again Blames ‘Both Sides’. <https://www.nytimes.com/2017/08/15/us/politics/trump-press-conference-charlottesville.html>. (2017).
- [16] David Smith. 2017. Trump fires FBI director Comey, raising questions over Russia investigation. <https://www.theguardian.com/us-news/2017/may/09/james-comey-fbi-fired-donald-trump>. (2017).
- [17] David Smith and Sabrina Siddiqui. 2017. ‘I love it’: Donald Trump Jr posts emails from Russia offering material on Clinton. <https://www.theguardian.com/us-news/>

- 2017/jul/11/donald-trump-jr-email-chain-russia-hillary-clinton. (2017).
- [18] Peter Snyder, Periwinkle Doerfler, Chris Kanich, and Damon McCoy. 2017. Fifteen minutes of unwanted fame: detecting and characterizing doxing. In *IMC*.
- [19] Andrew Torba. 2016. How You Can Help Support Gab. <https://medium.com/@Torbahax/gab-donations-9ca2a5c0557e>. (2016).
- [20] Mike Wendling. 2016. The saga of 'Pizzagate': The fake story that shows how conspiracy theories spread. <http://www.bbc.com/news/blogs-trending-38156985>. (2016).
- [21] Wikipedia. 1964. Jacobellis v. Ohio. https://en.wikipedia.org/wiki/Jacobellis_v._Ohio. (1964).
- [22] Wikipedia. 2017. Unite the Right rally. https://en.wikipedia.org/wiki/Unite_the_Right_rally. (2017).
- [23] Jason Wilson. 2016. Gab: alt-right's social media alternative attracts users banned from Twitter. <https://www.theguardian.com/media/2016/nov/17/gab-alt-right-social-media-twitter>. (2016).
- [24] Savvas Zannettou, Tristan Caulfield, Emiliano De Cristofaro, Nicolas Kourtellis, Ilias Leontiadis, Michael Sirivianos, Gianluca Stringhini, and Jeremy Blackburn. 2017. The Web Centipede: Understanding How Web Communities Influence Each Other Through the Lens of Mainstream and Alternative News Sources. In *IMC*.